

## DESCRIPTION

## MULTI-TONE TRANSMISSION

5 Field of Invention

The invention relates to multi-tone transmission, particularly for transmitting digital data.

Two commonly assigned patent applications in the name Mark Taunton, filed on the same date as the present application, each also entitled "Multi-tone  
10 transmission", and with reference numbers BP1756 US and BP1757 US are incorporated herein by reference.

Background Art

A system using multiple tone signalling generally uses the Fourier  
15 Transform and its inverse to convert the information between time and frequency domains. Two examples of this type of modulation scheme are: (a) DMT (Discrete Multi-Tone) as used in systems such as ADSL (Asymmetric Digital Subscriber Loop); and (b) COFDM (Carrierless Orthogonal Frequency Division Multiplex), a standard widely adopted for digital terrestrial TV  
20 broadcasting.

In these systems, the data to be transmitted are sub-divided (multiplexed) across a number of distinct frequencies (sometimes also referred to as tones or sub-carriers) which are all integer multiples of a fixed basic frequency. The individual tones making up the group are spaced apart by this  
25 basic frequency. (In the case of COFDM the group of tones is then shifted up to a much higher frequency range for transmission from an aerial, but that detail is not relevant to the discussion here.) The number of tones used in different systems and within an individual system can vary, anywhere from 10 or so; e.g. for a low bandwidth ADSL upstream link, up to several thousand,  
30 e.g. an "8K-carrier" COFDM digital TV transmission.

The key algorithm common to the communication systems under consideration is the Fourier Transform, a mathematical scheme in which a

time-varying signal is represented not as a set of values in time but as the sum of a set of sinusoidal waveforms. Each sinusoid in the set has a distinct frequency which is an integer multiple of a base frequency called the analysis frequency. Fourier Transform theory shows that any varying signal can be  
 5 alternately represented in this way, by defining the unique set of amplitude and phase values for the individual sinusoids which sum together to form the signal wave-shape.

In the general (continuous) case, the size of the set of sinusoids is infinite and the spacing of the individual frequencies is infinitesimal. However  
 10 the particular type of Fourier Transform used in practical communications systems is the Discrete Fourier Transform (DFT). The term 'discrete' is used because the data is processed as a set of distinct samples, not a continuous signal. When a finite sequence of samples is transformed in this way, the size of the set of sinusoids that represent the signal in the frequency domain is also  
 15 finite. Hereafter, when the Fourier Transform is mentioned, the term 'discrete' should be assumed.

In summary, the normal ("forward") Fourier Transform is used to convert from a series of samples taken in the time domain into an equivalent representation of the same information, namely as a series of values in the  
 20 frequency domain, describing the amplitude and phase of each of a set of harmonically related sinusoidal waveforms. The reverse process, the Inverse Fourier Transform, performs the opposite operation, summing the waveforms described by the individual amplitude and phase values to re-create a composite waveform as a series of samples in the time domain.

25 The Fourier Transform and its inverse are relatively complex functions, but they may be implemented without difficulty using well-known algorithms on a digital signal processor. In particular, highly efficient versions of the transforms are known, commonly called the Fast Fourier Transform (FFT) and the Inverse FFT or IFFT, which operate on sample sequences whose lengths  
 30 are powers of 2, e.g. 256 points or 512 points.

The FFT and IFFT together provide for efficient encoding and decoding of signals. In a transmitter, a set of data bits may be encoded by the IFFT,

choosing particular combinations of amplitude and phase for each of the constituent frequency components to represent different data values. After all the data is encoded into the amplitude and phase of each constituent tone, the IFFT is performed to create a time-domain signal which is then transmitted.

5 For example, it is possible to encode 2 bits of data, representing 4 different possible values (00,01,10,11), on to one tone by simple quadrature modulation, where the amplitude is held constant and four distinct phase values (e.g. +45,+135,+225,+315 degrees, i.e. 90 degrees apart) represent the 4 different combinations. More complex mappings are possible (allowing more  
10 bits to be encoded on one tone) using more phase values, or combinations of different amplitudes as well as phases. In practical systems, modulation of one tone can be varied so as to represent as many as 15 or 16 bits in the best case (using 32768 or 65536 distinct combinations of amplitude and phase). Therefore in systems using hundreds of tones, some thousands of bits can be  
15 carried in each symbol in good circumstances.

The (forward) FFT is used at the receiver to reverse the process. Once time synchronization with the transmitted waveform has been achieved and equalisation for frequency-dependent phase and amplitude changes (inevitable in the transfer of the signal from transmitter to receiver) has been  
20 performed, the FFT is applied to the set of samples making up each received symbol, to reconstruct values of amplitude and phase for each of the tones in use. In general the values obtained by this process are not exactly the same as were initially encoded, for various reasons, including particularly the presence of noise introduced along the transmission path of the signal. Noise  
25 is unavoidable in any practical system. However, by applying various techniques to compensate for errors caused by noise, the original data may be recovered with an acceptable level of reliability, provided the system has been configured appropriately, taking into account the signal-transfer characteristics of the transmission path.

30 In order to ease the work of the receiver in equalizing the received signal for the effects of the transmission route, it is common to insert a short delay between consecutive symbols transmitted. In ADSL, this delay period is

called the "cyclic prefix time", in which what is transmitted is a portion of the signal extracted from the end of the immediately following symbol. The name "cyclic prefix" time derives from the fact that the short sequence has been used as a prefix to the new symbol and is cyclically congruent with it. Note that after  
 5 equalization, the signal received during the cyclic prefix time is ignored by the receiver. In COFDM, the delay period is called the "guard time"; no signal is transmitted during this time.

The IFFT-FFT (encoding-decoding) process provides for great flexibility in the communications system. Different frequencies in the spectrum covered  
 10 by the set of tones may have different characteristics in respect of noisiness and attenuation over the communication link (e.g. the phone line in the case of an ADSL system). By varying the encoding details tone by tone, this may be accounted for, so as to maximize the number of bits carried by the symbol in total, even when a particular single tone can only carry a small number of bits.  
 15 U.S. Patent No. 4,679,227, which describes multi-tone encoding schemes, presents techniques for accomplishing this.

One property of this type of signal encoding is particularly relevant. The waveform resulting from the IFFT can in principle have very large peak values in it - relative to the average amplitude of the signal as a whole - at points  
 20 where the particular phases of the individual tones happen to sum together in the same direction. For example, if all tones were using encoded simple 2-bit quadrature modulation, and all the data bits being modulated were zero (or more generally if the same pair of bit values were modulating each tone), then at the start of the time domain symbol created by the IFFT there would be a  
 25 high amplitude "spike", since each component waveform would have a real positive value 0.707 times its peak amplitude, and these would all sum together in the same direction. By contrast, if there is a general haphazard distribution of 1's and 0's in the data, the expected peak value in the average symbol would be much lower, although once in a while peaks will still occur.

30 On observing the output from a sequence of IFFT operations used to encode a (generalised) data sequence for transmission, the signal is seen to have a sample amplitude distribution which is very like random noise, when

considered on a statistical basis. The most frequently occurring sample amplitudes are those near zero (the central point – the distribution is symmetrical either side of zero). Higher amplitudes are less likely, but still occur, and there is a gradual reduction in likelihood of occurrence with increasing amplitude. The very highest sample amplitudes which can occur - unlike with true noise there is a finite limit because we use a discrete IFFT over a finite number of tones - are still many times higher than the average signal amplitude; however, such values occur only extremely infrequently.

The overall statistical properties of the sequence are complex. However, one simple measure of the properties of signals generally is their crest factor. The crest factor of a repetitive signal is defined as the ratio of its peak amplitude to its average (RMS) amplitude. Different types of waveforms can have very different crest factors, depending on their shape. For example a simple pulse waveform, where the signal jumps between just two levels +A and -A, has a crest factor of 1, i.e. the average and peak levels of the signal are the same. A simple continuous sine wave has a crest factor of  $\sqrt{2}$  (1.4142135...). Other wave shapes can be envisaged having widely differing crest factors.

When we are dealing with irregular (non-repeating) signals, such as the output from an IFFT process applied to a random stream of data, the definition of crest factor is adjusted. This is necessary, in order to take into account the statistical spread of amplitude values. In such cases we define the effective crest factor to be the ratio of a threshold level to the average (RMS) level of the signal overall, where the threshold level is that which only some particular small fraction (e.g.  $1/10,000,000^{\text{th}}$ , or  $10^{-7}$ ) of the generated samples will equal or exceed.

With signals created by an IFFT-based modulator, in general, systems in which few tones are used will have a smaller effective crest factor than systems with large numbers of tones. In a typical ADSL system, using 220 tones on the downstream path, the effective crest factor is around 5.3 at the  $10^{-7}$  probability threshold.

In practical systems based on the IFFT/FFT pairing, various steps are taken to reduce the impact of its sensitivity to regular patterns of input data. These can readily occur in data sequences delivered to an encoder, especially in the case of ADSL where a fixed padding data pattern must be inserted when  
5 no user data is waiting to be transmitted. The problem of such regular patterns in the original data causing high peaks in the output of the IFFT is usually dealt with by performing a reversible "scrambling" operation on the data stream prior to encoding. Two examples of such scrambling mechanisms are self-scramblers and randomisers.

10 By applying scrambling processes to the input data, any regular patterns in it may be broken up. The distribution of the data bit values going forward into the encoder becomes more haphazard, and so the likelihood of coherence between the phases of the different tones is drastically reduced. This diminishes the frequency with which spikes appear in the time-domain  
15 signal, even for a completely regular input stream (e.g. all 1s), relative to that which would apply without scrambling. However, for more irregular input data, no particular change in the statistical properties of the IFFT output will occur.

One major problem with IFFT-based encoding, so far as the design of any practical system is concerned, is that the time domain signal created has  
20 characteristics which make it more difficult and/or more expensive to carry through the later stages of the transmission path. For example, the bandwidth of the signal may in some cases be as wide as can theoretically be carried by the discrete sample sequence. Any subsequent processing of the signal, post-  
25 IFFT, must therefore be carefully designed to minimise distortions of the signal caused by frequency-dependent variations (e.g. in gain or phase-shift), which are typically worst at the highest frequencies.

However, an issue of great concern is the high crest factor of a typical IFFT-generated signal. This leads to a number of difficulties in designing the circuitry in a modulator & transmitter for an IFFT-based modulation scheme.  
30 Some of the problems also occur in the design of a corresponding multi-tone receiving device.

The first problem is that the dynamic range of the digital-to-analogue converter (DAC) must be large, requiring a relatively high number of bits of resolution (typically between 14 and 16 for ADSL). This makes the DAC hard to design, especially since it is running at high sampling rates (in the order of 5 1-2 MHz or higher for ADSL, and higher still for COFDM). In a receiver for the transmitted signal, the input circuitry must also have a high dynamic range and low noise and distortion; equally its analogue-to-digital converter must have high linearity and resolution.

The second aspect, which is usually considered even more serious, is 10 that it is extremely difficult to design the amplification stages of the transmitter to both yield the high linearity which is needed and also maintain good power efficiency. Because the amplifier (also called the "line-driver" in the case of ADSL) must be able to handle signal peaks several times higher than the average signal level on the line, it becomes necessary to run its power supply 15 at a far higher voltage than the average signal level would require, if the signal's crest factor were lower. Typical power efficiencies for amplifiers in present-day ADSL system designs are therefore significantly lower than in some other types of transmission system e.g. 15-20% as against 40% or more.

Accordingly, it would be beneficial to reduce the crest factor.

20 An existing patent which describes a technique for reducing the crest factor is US Patent No. 5,768,318. In that design, the peak value in each symbol produced by the IFFT is first found, by checking each sample which it contains. If the peak lies below the defined threshold, nothing further need be done and the symbol can be transmitted. Otherwise, when the threshold has 25 been exceeded, a special modification is applied to the input vectors of the IFFT (representing the phase and magnitude of each tone), and the IFFT operation is re-run on the modified vectors. The modification is defined in terms of a scaling and phase rotation for each vector. In the general case described in the prior art, this operation may be applied repeatedly, using 30 different modifications of the vectors each time, selecting whichever first results in a peak below the threshold in the resultant time-domain sample sequence created for the symbol.

The prior art patent also describes a possible implementation in which multiple modifications of the vectors and the IFFT operations on each set of modified vectors are executed in parallel, with a selection being performed between the different resulting time-domain waveforms (each held in a buffer) for one which gives a suitable low peak. This scheme reduces the time penalty implied by repeated sequential vector-modification and IFFT operations, but at the cost of greater hardware size, including more buffers.

Now in order for the decoder at the receiving equipment to know what modification of the vectors (if any) was performed, and hence be able to reverse it so as to decode the original information carried, two things are necessary: (a) the modifications are done in accordance with a fixed set of definitions known by both transmitter and receiver; (b) at least one tone must be reserved in the signal spectrum used to carry the symbol. The data modulated onto the reserved tone(s) is an index value ranging over  $[0..N]$  identifying whether any modification was performed, and if so which out of  $N$  possible modifications was actually used in the symbol as transmitted. The reserved tone(s) clearly must not be modified themselves, and are not available to carry ordinary data.

The main disadvantages of this scheme are:

(a) It requires both transmitter and receiver to co-operate in order to apply the process. The receiver is required to agree with the transmitter about the reserved tones used to carry the modification index. Then, on reception of every symbol, the receiver must first identify the index of the modification applied (by decoding the reserved tones), and if appropriate perform the reverse processing on the output of its FFT operation, before it can continue with the decoding process. This will take extra computation capability, and time, at the receiver.

(b) It consumes data bandwidth for the reserved tone(s) to carry the modification index information; this reduces the bandwidth left for user data.

(c) The carriage of the modification index value must be especially reliable - if the index is decoded incorrectly, the whole symbol cannot be decoded. To provide the additional reliability, either extra bits for forward error



correction (such as a CRC), must be sent along with the index value itself, or else the degree of noise margin required when modulating the reserved tone(s) must be set much higher than normal for the main data-carrying tones, or some other method of protection must be applied to that information.

- 5 Whatever scheme is used will inevitably take more data bandwidth than the same number of bits would need in the main body of user data, so the available user data bandwidth is further reduced.

(d) In the context of standards-conformant ADSL systems, the technique is incompatible with existing standards. Only if the two modems  
10 involved both support the technique in the same way can it be used. To do so they must use standards-compatible methods to identify each other's capability in this regard, early in the negotiation phase of initialising the link, before the tones have been allocated to data. Then some proprietary method must be used to set up and activate the extra processing involved at each end of the  
15 link. There is no possibility of using this technique when only one of the modems supports it, or if they use differing proprietary methods to implement it. Hence its application is limited.

Accordingly, there remains a need for an improved method of crest factor reduction and improved apparatus implementing the improved method.

20

#### Summary of Invention

According to a first aspect of the invention there is provided a modulation method for multiple-tone signalling, including the steps of:  
processing input data through a plurality of intermediate processing stages  
25 and corresponding stages of intermediate data to generate preprocessed data; inverse Fourier transforming the preprocessed data to obtain a symbol including a number of tones; buffering the symbol; determining the peak amplitude in the symbol and comparing the peak amplitude with a threshold; if the peak amplitude in the buffered symbol exceeds the threshold, obtaining a  
30 regenerated symbol by amending predetermined intermediate data such that the input data is still represented by the intermediate data, carrying out the subsequent intermediate processing stages on the amended intermediate data

to regenerate preprocessed data, and inverse Fourier transforming the regenerated preprocessed data to obtain the regenerated symbol, replacing the buffered symbol with the regenerated symbol; and outputting the buffered symbol.

5        It will be appreciated that the method of the invention will generally be applied to a stream of incoming data to output a stream of symbols. Normally, only a very small proportion of the symbols will need to be regenerated.

10        The idea underlying the invention to be described here is that, when there is a peak above the threshold in the time-domain symbol created for transmission, re-processing to create a new and different time-domain version of the symbol is performed, not directly on the input vectors of the IFFT, but starting further back in the sequence of operations which lead to the generation of the IFFT input vectors. Furthermore, instead of performing some special reversible modification of the intermediate values representing the  
15        symbol (such as the phase and magnitude information of the IFFT input vectors), leaving the data carried unmodified, and having to send additional information to the receiver to indicate when this has been done, it is possible to change certain data bits being carried within the symbol as transmitted, but in such a way as to be of no consequence for the receiver, and not to affect  
20        the reliability of transmission of the user data.

      In another aspect there is provided a modulation method for multiple-tone signalling, including the steps of: processing input data through a plurality of intermediate processing stages and corresponding stages of intermediate data to generate preprocessed data; transforming the preprocessed data to  
25        obtain a symbol including a number of tones; and if the peak amplitude in the symbol exceeds a predetermined threshold, regenerating the symbol by amending predetermined intermediate data such that the input data is still represented by the intermediate data, and carrying out the subsequent intermediate processing stages and the transforming step on the amended  
30        intermediate data.

      The invention relates to apparatus for carrying out the methods as well as to a computer program product for implementing the method.

Using the invention it is not necessary to introduce extra complexity by defining new, relatively expensive special operations on the IFFT input vectors, as the prior art does, in order to achieve the goal. Nor is it necessary to pass a separate indication along with the main data in the symbol, to allow the receiver to reverse the process and recover the user data. Instead all that happens is that a modification is made to some intermediate data item in the set of data items which were encoded to create the symbol, followed by a repeat of the normal data processing for transmission, using exactly the same steps as were performed the first time the symbol was encoded (and which any transmitting modulator would perform). Thus the receiver can be left completely unaware that the transmitting modem is implementing this scheme; it has nothing different to do. The same amount of user data can be transmitted, in exactly the same way as before, and no extra information has to be transmitted. Thus, the disadvantages inherent in the prior art arrangement are avoided.

The key discovery giving rise to this invention is that significant perturbations in the values of the input vectors for the IFFT, sufficient to cause the modified symbol to take a shape substantially different from the original one, can be accomplished by quite small and simple changes to the data being processed within the encoding system prior to the IFFT input stage. When a symbol is detected whose final (time-domain) signal shape contains a peak higher than the threshold level, all or part of the processing of data which was performed in order to create that symbol is re-executed, this time making a change to some item of data which contributes to the symbol. If the resulting re-generated symbol has a lower peak value than the threshold, all is well and the revised symbol is sent. Depending on circumstances, (e.g. as determined by the availability of required resources such as processing cycles and buffer memory) if the peak value of the new symbol is still no lower than it was, or still above the threshold, further attempts at re-generation may be made, using other changes in the data.

The one known disadvantage of this new scheme relative to the prior art, is that a little more computational cost is incurred, in that when it is desired

to re-generate a symbol, some of the earlier processing prior to the IFFT has to be re-run, as well as the IFFT operation itself; however in compensation no separate processing of the IFFT input vectors is required. Such re-processing may also require some additional memory in order to buffer data in the modem  
5 for a longer period than it would otherwise be needed for. However, the overall cost of this will typically be small, especially for a system in which the repeated processing is performed in software rather than by fixed hardware functional blocks (though a hardware-based system could also be designed to implement this invention), and even more so when a single device is implementing the  
10 modulation processing for multiple channels of ADSL transmission at the same time; then the small extra resources required can be shared between all instances of the modulator function – the frequency with which they are needed will typically be low enough that only one channel is likely to need them at any one time.

15 By way of illustration, if it is desired to reduce the crest factor of a transmitted ADSL signal from 5.3 to 4.3, and assuming that any symbol may be freely regenerated, on average only one symbol in 130 will need to be re-generated for any given line; for a target crest factor of 4.0, the average regeneration rate is approximately one in 25 symbols.

20 Preferably, the subsequent intermediate preprocessing stages used to regenerate preprocessed data include a scrambling stage. A scrambling stage is efficient at causing significant changes in the final symbol for only small changes in the intermediate data.

In embodiments, the data is structured as a series of frames including  
25 data bytes from an ATM cell stream (or alternatively, the data bytes may come from an STM data byte stream) and in at least some frames further including fast bytes and/or sync bytes and the step of amending the predetermined intermediate data includes amending the fast or sync bytes of the frame sequence.

30

#### Brief Description of the Drawings

Embodiments of the invention will now be described, purely by way of example, with reference to the accompanying drawings, in which.

Figure 1 shows a block diagram of a first embodiment of the invention;

Figure 2 shows a block flow diagram of a second embodiment of the invention; and

Figure 3 shows a block diagram of a transmitter implementing the second embodiment.

#### Specific Description

Figure 1 shows a schematic diagram of a first embodiment of an apparatus according to the invention.

An input data stream 10 is input through a plurality of preprocessing stages 12, 14, 16. Each preprocessing stage produces intermediate data 18, 20, 22. The intermediate data 22 produced by the final preprocessing stage 16 will be referred to as the preprocessed data 22.

The skilled person will be aware of many suitable intermediate preprocessing stages that can be used or not used as required. For example, the input data may need to be divided into cells or frames. Buffering may be required. Blank data can be added if insufficient user data is available. One or more of the stages may carry out scrambling, so that if the input data is regular the output of the scrambling stage is not regular to avoid peaks in the output. Error correction, such as Reed-Solomon, and other suitable steps may be used.

Further intermediate preprocessing stages may be used to generate the multi-tone symbols. One stage may assign the data to tones, another may carry out constellation encoding and a further stage may carry out gain scaling. The skilled person will be able to select suitable stages for the required application.

The preprocessed data 22 is fed into the Fourier transform module 24 that carries out an inverse discrete Fourier transform to generate a symbol 26, which is buffered in a symbol buffer 158.

The peak amplitude in the symbol is determined in a peak amplitude detector 28. This detection is preferably carried out in the digital domain immediately after the IFFT is carried out, though it may be possible in some embodiments to carry out the determination at a later stage.

5 If the peak amplitude exceeds a predetermined threshold, then the symbol is regenerated. Intermediate data 18 input into one of the intermediate stages is amended, or regenerated, in such a way that the regenerated data still corresponds to the original input data, by changing one or more bits in the intermediate data. The regenerated intermediate data is then processed  
10 through subsequent intermediate stages 14, 16, and the Fourier transform module 26 to produce a regenerated symbol.

Conveniently, an intermediate data buffer 30 may be provided in one or more of the intermediate stages to retain the intermediate data 18 in case it is required to regenerate the data.

15 There are a number of possibilities as to what occurs with the regenerated symbol. The simplest option is simply to replace the symbol in the symbol buffer 158 with the regenerated symbol. Since the original symbol in the symbol buffer exceeds the peak threshold, replacing the symbol is likely to reduce the peak amplitude and hence improve the crest factor, assuming  
20 the threshold for regeneration is correctly set.

In an alternative arrangement, the regenerated symbol may be stored in the symbol buffer 158 as well as the original symbol. Then, the peak values of the original and regenerated symbol are compared, and the original symbol replaced if the regenerated symbol has a lower peak value. The replacement  
25 may take place in any convenient fashion, including for example copying the regenerated symbol over the original symbol or simply swapping the functions of the two stored symbols.

If the peak amplitude remains too high a further regeneration attempt, or attempts may be made. The best symbol, with the lowest peak amplitude,  
30 may finally be sent. In embodiments, up to a predetermined maximum number of attempts may be made; the predetermined maximum number may be in the range three to ten inclusive, for example six.

In some cases, there may not be sufficient processing resources or time to make a regeneration attempt or attempts. Therefore, it may be determined whether or not there is sufficient time before regenerating a symbol. As will be appreciated, it is desirable not to provide excessive processing capacity.

5 Although the diagram shows a plurality of separate stages, some or all of the stages may be implemented using a multi-purpose data processing apparatus programmed to carry out the same method.

The main discovery embodied in this invention is that various of the stages of processing performed on the stream of data passing through the modulator in a multi-tone modem (from its data input towards the physical line transmitter) have a general tendency to "spread" the original data values over a wider distance in the later stages of the stream than they occupy in the input. (This is true for both ADSL and other types of modulation system using multi-tone signalling.) By this is meant that a change in a single bit or group of bits  
10 in the input can cause a larger number of bits to be changed in the subsequent output stream from some data processing steps. Thus the scale of change in the data processing, required to cause a large perturbation in the input vectors of the IFFT (i.e. one large enough to mean that the re-generated symbol is effectively a new and independent shape) can be kept very low, potentially as  
15 low as a single bit changing.

In the specific embodiment, the step 14 following the introduction of regenerated intermediate data 18 is a scrambling step. Such a step has a particularly strong effect: any single bit change (0 to 1 or 1 to 0) in the payload of the ATM cell stream will have a permanent effect on the subsequent output  
20 from the payload scrambler (after an initial period while data bits percolate through its internal memory), since that process has local feedback and retains a memory of the effect of all bits which have passed through. The second scrambling step, applied to the full data stream (including ATM cell headers and other overhead data), similarly causes all later data output values (beyond  
25 an initial short sequence) to be affected as a result of changing a single input bit.

The scrambling step used in the embodiment is a self scrambling step, i.e. a scrambling step using an internal memory of previous outputs or previous states which are used together with new data to generate new outputs in a recursive manner. This provides a useful "spreading" effect allowing small changes in input data to produce large changes in the output symbol.

It is however not necessary that the subsequent steps include a self-scrambling step, or indeed a scrambling step at all. Other steps, such as Reed-Solomon or Trellis encoding also produce suitable spreading. Furthermore, the scrambling step or other spreading steps need not be carried out immediately on the amended intermediate data, but the steps doing significant randomisation may also be subsequent steps.

Moreover, other operations such as the addition of Reed-Solomon error protection bytes, and the application of trellis encoding, have a weaker effect, but one which is still very helpful for present purposes. For Reed-Solomon (R-S) protection, changing any one bit in a block of data bytes which is protected by R-S parity bytes will generally cause many bits of those parity bytes to change (but will have no other effect on the rest of the data block). The spreading effect with trellis encoding is most pronounced for tones which have few bits allocated. Since in common usage of ADSL, many tones do indeed have few bits (e.g. 2, 3 or 4 bits per tone), this is still a significant effect.

The present invention may be applied to ADSL modems. Use is made of this "spreading" effect of the various data processing stages, so as to cause a large change in the output of the final processing step before the IFFT is performed, thereby yielding a different shape and hence a different and probably lower peak value of the time-domain symbol. To do so it is required to introduce only a small change to the data passing through an early step of the processing sequence. Furthermore as will be shown, such changes can be made without impacting correct transmission of the incoming data stream through the communication path of which the ADSL modem forms one part.

Alternatively, the invention may be applied to other types of multi-tone signalling systems, such as COFDM (Carrierless Orthogonal Frequency



Division Multiplex), a standard widely adopted for digital terrestrial TV broadcasting.

In the case of digital terrestrial TV broadcasting using COFDM, the scrambling process is different from that of ADSL, being of a randomising type rather than self-scrambling, and causes little or no spreading of data bit changes. However the trellis/convolutional coding process used for COFDM is very effective in this respect - much more so than ADSL trellis coding. Thus, for COFDM the trellis and convolutional coding may be used for spreading the data.

A second embodiment will now be described, specifically with reference to an ADSL system. Figure 2 shows a flow diagram of an ADSL modem that transmits ATM cells 102. In a first data path 100, the cells are buffered 104, and idle cells 106 are inserted 108 as required. The cell payload is then scrambled 110, and a cyclic redundancy check 112 performed. Data bytes from the ATM cell stream are then grouped together by framing 114, adding fast bytes 116 where required. Scrambling 118 then takes place, followed optionally by Reed-Solomon forward error correction 120.

Another, interleaved data path 124 is also shown, having the same steps except that sync bytes 117 are used instead of fast bytes 116, and there is additionally a final step of convolutional interleaving 125.

The cells of the two data paths 100, 124 can then be merged 122. As will be appreciated, it is not essential for there to be any particular number of data paths and the merging step 122 is only required where there are a plurality of paths.

The framed, scrambled and merged data stream is then passed to the modulator 126 which carries out the steps of tone ordering 128, optional trellis encoding 130, constellation encoding 132, gain scaling 134, and finally inverse Fourier transform 136 to produce a stream of symbols each encoding some portion of the ATM cell stream, ranging from one or two partial ATM cells up to several whole and some partial ATM cells. Until this point, all of the operations are conventional and well known to those skilled in the art.

The peak amplitude in the symbols produced by the modulator is then detected 138 and compared with a predetermined value.

If the peak is above the predetermined value then regeneration of the symbol is performed, under regeneration control 140. The ways in which this is done will be described in more detail later.

Normally, however, the symbols from the modulator are simply buffered 142. Then, on transmission cyclic prefixes are inserted 144 before the data stream is passed to the analogue front end 146.

In the analogue front end, digital to analogue conversion 148 is performed following which the stream is passed to a line driver to drive 152 the telephone line 154.

Note that each functional block shown in the Figure 2 within the Data Path modules, the Modulator module and the Regenerator Control module, could in principle be implemented either by hardware or by software, or by some combination of the two. The AFE module normally uses hardware blocks for its functions.

A specific embodiment of an implementation is illustrated in Figure 3. A modulator 126 feeds data through a buffer 158 into the analogue front end 146 which contains a DAC 156 and a line driver 150 to drive the phone line 154. These features are known to the skilled person and will not be described further.

The symbol peak detection & regeneration control logic is in the modulator 126 as shown in Figure 2. This logic is an important part of this invention. Note that the output from the regeneration control logic block is not shown explicitly - it signals that re-generation should be performed for a given symbol; this signal must be conveyed to a number of places in the rest of the modem, to trigger the re-generation process. In a modem implemented (at least partly) in software, this indication would cause a re-execution of the relevant parts of the illustrated modem functionality performed by the software code.

To implement the invention in an ADSL modem, certain features are provided in the device. For a number of known modem designs, some of

these features are already present, so use of the invention may require changes of only quite modest scope, compared with an existing design.

One feature is the buffer 158 between the output of the IFFT processing stage and subsequent stages where the generated symbol data are passed on  
5 towards the output circuitry of the modem. This buffer is of sufficient size to hold at least all the samples representing one full time-domain symbol, plus additional sample buffering to cover the transfer of data from earlier output symbols while symbol-regeneration is attempted. The buffering is used because when a symbol is found after IFFT to contain a peak which is above  
10 the threshold, the symbol must not be allowed to proceed further, into the final output stages of the modem, but must be held pending an attempt or attempts to re-generate it. If re-generation is successful, the new symbol created must replace the original one in the buffer before the latter is ready to be transmitted down the line; failing that, the original symbol may be sent.

15 Some additional buffering may also be required earlier in the data processing parts of the modem, in order to support the re-processing capability. Specifically, any data which would normally be discarded once the data for a particular symbol had been processed through a stage, must instead be retained in case re-processing of that symbol is required. However, this  
20 additional memory is only needed until the check is made as to whether the time domain symbol contains a peak above the threshold, prompting a re-generation attempt. The memory may be re-used as soon as a symbol's peak value is found to be below the threshold, implying that no re-generation is needed.

25 Preferably there is a reserve of processing capacity, over and above the amount required to perform normal modem processing for transmission. This is because when symbol re-generation occurs, additional processing time will be spent on this, which must be made up for, over time. Normal per-symbol processing should therefore take at least somewhat less than the real time for  
30 a symbol to be transmitted. For example, if it is expected that one in every 1000 symbols will need to be re-processed for purposes of peak reduction, then the individual per-symbol transmit processing cost needs to be less than

1-1/1000 = 99.9% of the available processing resource. Clearly this is not in general a difficult constraint to meet(!). However, to allow for irregular timing of re-generation events (such that while the average rate is 1 every 1000 symbols, sometimes two high-peak symbols may occur in close succession), requires a further margin of processing performance: e.g. if there is a 5% reserve (processing for each symbol is completed in only 95% of available time) then two difficult symbols only 20 symbol times apart can be handled. In general, statistical analysis based on the exact peak-reduction factor which is aimed at will allow a characterisation of the actual performance margin needed in a particular system design. Note that it is perfectly acceptable to have a rule which prevents overloading of the processor by refusing to attempt regeneration of a symbol if insufficient time has elapsed since the last re-generation event occurred (i.e. the processing has not "caught up" enough at the time from the earlier events). All that happens is that that particular symbol will be transmitted unmodified, having a peak above the threshold, but the overall (long term) crest factor will still be lower, because such situations will statistically be very rare relative to the frequency of effective symbol re-generation.

The actions carried out to regenerate symbols which exceed the predetermined peak value will now be described in much more detail. Various methods of changing the data values in the ADSL stream are identified, which can be applied in manner fully compatible with the various national and international standards for ADSL systems (e.g. ITU G.992.1). Note that these methods are not mutually exclusive, but can be applied in various combinations to allow multiple attempts (either sequentially or in parallel) when re-generating a symbol.

#### Fast byte/sync byte "X" bits

The single most valuable point at which data changes can readily be performed is in the "fast" and/or "sync" bytes which are defined to occur in ADSL data symbols. The fast and sync bytes are overhead bytes, not

themselves part of the stream of data (usually an ATM cell stream) to be carried over the link, but associated with it, and physically carried as a part of the modulated signal. Depending on configuration, in many (though not all) ADSL data symbols, they contain control information used to manage the synchronization of data streams being transported over the ADSL link which were originated via a communication path whose control clock is asynchronous to the ADSL modem's own control clock. However in many practical systems using ADSL modems, this capability is not required. Even where it is needed, it may actually be applied only rarely, leaving the byte available for application of the technique described here, most of the time.

When a fast or sync byte is not carrying synchronization control values, it is defined to carry values of a fixed pattern, of the form XX0011X0 for the fast byte and XX0011XX for the sync byte. Each bit shown as X can be freely set to either 0 or 1 as desired. With three or four bits whose value may be changed at will, there are a total of 8 or 16 possible combinations of 1s and 0s which may be created; therefore up to 7 or 15 attempts to re-generate a symbol are possible. This is more than adequate, in general.

Depending on configuration, fast bytes occur at the start of many (e.g. 64 out of every 68) data frames of ADSL when "fast" (low latency) data streams are used. Changes in any X bit of the fast byte will cause much larger scale changes in the symbol because the stream scrambler is applied to this byte first in the fast data stream, so the scrambled form of almost all subsequent data bytes in the fast stream (but not the interleaved stream if also present) will in general be altered. Reed-Solomon encoding (if applied) also follows the scrambler so the added R-S parity bytes will in general take different values. Finally, the trellis encoder (if applied) will also cause changes to the output stream, because of any change in its input, over a given frame. Furthermore, since the fast byte is the first byte in each whole data frame (including also the interleaved data if present), this means that the trellis encoding of all or almost all tones in the symbol is liable to be modified by a single bit change in the fast byte.

Sync bytes occur within most frames when interleaved streams are used. As with the fast byte, changes in any X bit of the normal sync byte pattern will affect all subsequent bytes in the interleaved portion of the data frame, through the application of the stream scrambler to the interleaved stream. However if the convolutional interleaver is applied, then older data from the interleaver's buffer, which will also appear in the final encoded symbol, will not be affected by the change; thus changes to the sync bytes are less effective. However, if trellis encoding is in use, then all output data of the trellis encoder, starting from the interleaved part of the frame, will still be affected by a change in the sync byte.

Both fast and sync bytes occur in many frames in a "dual latency" system where both fast and interleaved streams are active. In this case, either or both fast and sync bytes may be changed (a total of up to 7 X bits) to effect a significant change in the final time domain form of the encoded symbol.

The principal limitation of this method is that in some cases, the fast or sync bytes are not available - 4 data frames (symbols) out of the 68 data frames in each ADSL "superframe" structure are defined not to carry them (the byte location in the frame is used for a different purpose in these frames), and in systems where synchronization must be performed, these bytes may occasionally carry values other than the default pattern with its three X bits. In most configurations, the fast and sync byte locations in the ADSL data frame are shared with use for other purposes, though these will generally be infrequently used; additionally one configuration reduces the number of frames carrying fast and /or sync bytes to 32 (rather than 64) out of every 68. If it is desired to modify the data and regenerate the symbol in these cases also, some other method must be found; alternative methods are given below.

#### Idle cell payload modification

Of all possible techniques for changing data in an ADSL modem in order to cause symbol re-generation, changing data in the payload bytes of an ATM cell, as seen at the input end of the transmitter-side processing, has the greatest possible "spreading" effect. This is because the largest number of

steps, each of which causes spreading, will subsequently be applied to the data: (a) the payload bytes in the cell are scrambled; (b) the composite stream is then scrambled again; (c) Reed-Solomon error protection may optionally be applied to the stream; and (d) trellis encoding may optionally be applied to the stream while creating the IFFT input vectors. Even if no R-S encoding or trellis encoding is used, the two sequential scrambling processes are enough to cause a huge spreading effect on the data. Therefore, if even one bit of a cell's payload is changed, the regenerated data at later stages of the ADSL processing sequence will in general be very different, resulting in an equally significant change in the time-domain signal created, and hence (by probability) a high likelihood of creating a new symbol with a lower peak value than the original one.

However, in general, it is not desirable to change the user data – the purpose of the modem is faithfully to transmit exactly that data with which it was supplied, over the link to the other end of the ADSL connection. Therefore it is strongly preferred, not to modify that data. However, idle cells can provide a way to achieve the effect of the invention, without damage to real user data.

Idle cells are ATM cells of a special reserved type, which are used to pad out a data stream. Idle cells are defined by a particular fixed pattern in the cell header – this is how they are recognized as idle cells when received. The payload of an idle cell is also defined as a fixed pattern, the same in every byte. In ADSL, the transmitting section of each modem is obliged to insert idle cells into the data stream whenever no user data cells are available to be transported. This is required because the physical data rate of a standard ADSL link is fixed at initialisation, and is maintained until the line is shut down or re-initialised. Since it is not possible to send “no data” when no user data is present, idle cells are sent instead, to maintain the flow. At the receiving end, idle cells are simply discarded – their contents are not related to real user data carried by the connection.

This observation provides another method of changing the data stream, for symbols where it is desired to re-generate because of a peak above the

specified threshold in the encoded time-domain version of the symbol. What can be done is to check whether any byte of data carried by the symbol is part of the payload of an idle cell, and if so, to make a modification in that byte. Because the payload of idle cells is ignored by the receiver, any bit of the 8 in  
5 the byte can be modified, allowing up to 255 possible modifications (relative to the original value) to be tried out – this is far more than enough! In general, to achieve maximum effect, the earliest available idle-cell payload byte in the data for the symbol should be so modified, since all modifications affect (by spreading) only the encoded form of later bytes in the stream (and hence, that  
10 portion of the stream which is carried in the rest of the symbol).

Most installed ADSL connections only carry user data for a small proportion of the time, when considered on a long term basis; so idle cells will be very common in general. Since an idle cell does not contain any user data, its payload is not of interest, being fixed. Changing any bit in the payload of an  
15 idle cell will have no effect on the user data also carried by an ADSL connection.

This method is therefore applicable in many cases. Only when the line is fully occupied with user data at the time, such that no byte in the data carried in the particular symbol to be regenerated is part of an idle cell's  
20 payload, will this technique not work. But since on a long term statistical basis, most bytes of data carried over ADSL lines will be idle cell payload bytes, this method will frequently be available.

One proviso to this method is that one way of checking the error rate on an ADSL link, sometimes employed for purposes of link maintenance and  
25 management, is for the receiving modem to examine the payload bytes of idle cells before it discards them, comparing each byte against the fixed value it is defined to hold in any idle cell. Any errors found in the comparison are assumed to have arisen as a result of uncorrected errors in transmission of the data stream over the ADSL link. Some modems keep count of the error rate  
30 on this basis (measured as a moving average of the number of bits in idle cells which are found to be incorrect, divided by the total number of bits in the idle cells seen, over some measurement interval). In such cases, the deliberate



introduction of changes to idle cell payload bytes will give rise to an incorrect assessment of the true error rate by the receiving modem. In the worst case this may trigger an attempt to re-configure or re-initialise the link, so as to maintain the apparent error rate below the required maximum level.

5        Clearly it is useful to guard against such behaviour. One possible method is to define a limited set of modifications to idle cell payload bytes which can be attempted by the transmitting modem. Instead of the standard fixed payload byte value, a small number of alternative values (say, 3 out of the 255 remaining possibilities, or even just one value) could also be  
10       considered to be "legitimate" in idle cell payloads. In such a scheme, the receiving modem would be modified so as not to count such values in idle cell payloads as being errors, for purposes of error rate calculation. This modification still allows a high rate of true error detection, since the probability that a randomly corrupted idle cell payload byte takes one of 4 specific allowed  
15       values out of the 256 possible ones is only 1/64, or 1/256 if just one alternative value is allowed. Therefore with random, even error distribution, the true error rate and the measured one would differ by at most 6.25%, well within an appropriate level of accuracy in this context; furthermore in long term measurements it is possible to compensate for this difference.

20

#### Cell swapping

As an alternative to modifying the contents of (idle) ATM cells, another method to cause a change in the data stream so as to allow symbol re-generation is to physically replace one ATM cell by another. For the reasons  
25       already identified, this will in general give rise to a large change in the encoded time-domain form of the symbol. This type of technique can be applied to any symbol where the data it carries includes the start (first header byte) of an ATM cell. Various possibilities are available:

(a)    If the start of an idle cell was encoded in the original form of the  
30       symbol because no data cell was available for transmission as processing for that symbol was started, and if when the processing has been completed it is found that the encoded time-domain symbol should be re-generated, a check

should be made as to whether a new data cell has become available; if so then the idle cell can be discarded and replaced with the data cell.

(b) If the total ATM cell stream being carried over the link contains cells with different ATM addresses in their headers (which form logically distinct streams based on the addresses) and if at the time when symbol regeneration is required a new cell with a different address from that of a cell starting in the originally encoded symbol is also available for transmission, then the two cells can be swapped and the symbol regenerated using the second cell. This should not normally be done if the first cell was marked as being of higher priority than the first, or if the first cell had been waiting to be transmitted for much longer than the second.

(c) If the line is currently running at well below full capacity, (i.e. there is a strong likelihood of inserting an idle cell in the near future, based on the state of the ATM cell buffering at the input to the transmit side of the modem) and if the current symbol, containing the start of a user data cell, is to be re-generated, then an idle cell may be substituted for that data cell, and the data cell replaced at the head of the queue of cells awaiting transmission. This type of change should be attempted only if all other methods are inapplicable, since it causes a delay to the cell's transmission. In addition it should not be used if the cell concerned is marked as high priority.

As shown above, there are a number of viable ways of implementing the invention for ADSL, while maintaining correct operation of the system as a whole, without requiring the receiver to have any knowledge of the process at all, and in a completely standards-compatible manner. Additionally, very minor changes to the standards-defined processing (such as the allowance at the receiver of one or a small number of alternative legitimate values for the bytes of an idle cell payload) usefully extend the range of possible options for symbol regeneration. As mentioned above, the invention is not limited to ADSL.